

1N-14-CR
69294

An Investigation of Networking Techniques for the ASRM Facility

P-57

Final Report
February 4, 1992

Robert J. Moorhead II
Wayne D. Smith
Principal Investigators

Dale R. Thompson
Research Assistant

for the period
February 6, 1991 to February 5, 1992

Department of Electrical and Computer Engineering
Mississippi State University
Mississippi State, MS 39762
(601) 325-3912

Research Grant: NAG8-866

submitted to:
NASA/MSFC
EB32/Robinson

(NASA-CR-189851) AN INVESTIGATION OF
NETWORKING TECHNIQUES FOR THE ASRM FACILITY
Final Report, 6 Feb. 1991 - 5 Feb. 1992
(Mississippi State Univ.) 57 p CSCL 14B

N92-18286

Unclassified
G3/14 0069294

ABSTRACT

This report is based on the early design concepts for a communications network for the Advanced Solid Rocket Motor (ASRM) facility being built at Yellow Creek near Iuka, MS. The investigators have participated in the early design concepts and in the evaluation of the initial concepts.

The continuing system design effort and any modification of the plan will require a careful evaluation of the required bandwidth of the network, the capabilities of the protocol, and the requirements of the controllers and computers on the network. The overall network, which is heterogeneous in protocol and bandwidth, is being modeled, analyzed, simulated, and tested to obtain some degree of confidence in its performance capabilities and in its performance under nominal and heavy loads. The results of the proposed work should have an impact on the design and operation of the ASRM facility.

INTRODUCTION

In February 1991, the proposed Manufacturing network consisted of an FDDI ring off of which hung 5 Ethernet fiber-based LANs and the OIS computer. After this configuration was analyzed, the results were summarized, submitted, and subsequently accepted as a refereed paper to the IEEE Southeastcon '92 conference. A copy of the paper is included in Section 1. Several changes were made to the network as the year progressed. Once the OIS computer (a 4-machine VAXcluster) was procured, the FDDI ring disappeared and the 5 LANs were now attached directly to the OIS computer. This configuration was simulated and the results are discussed in Section 2. In July 1991, as the data rate requirements began to decrease, a data-over-voice network was proposed for non-critical sections of the network. The current proposed network consists of a hybrid of Ethernet-over-fiber and Intecom LANmark. This current proposed network is discussed in Section 3. Sections 1, 2, and 3 will each be presented as if they are the final solution; minimal attempt is made to reword the results based on design decisions that came later. The two network technologies (Ethernet over fiber and LANmark) were tested on November 12-13, 1991 and on December 12-13, 1991 in Iuka, MS at the ASRM facility and the results are discussed in Sections 4 and 5. Section 6 consists of a summary and some conclusions on where the network now stands and where the design seems to be headed.

SECTION 1

ADVANCED SOLID ROCKET MOTOR (ASRM) COMMUNICATIONS NETWORK ANALYSIS USING BONES

REFEREED PAPER SUBMITTED TO IEEE SOUTHEASTCON '92

ADVANCED SOLID ROCKET MOTOR (ASRM) COMMUNICATIONS NETWORK ANALYSIS USING BONeS

D. R. Thompson, R. J. Moorhead, and W. D. Smith

Department of Electrical and Computer Engineering

Mississippi State University

Mississippi State, MS 39762

ABSTRACT

This paper describes the simulation of a proposed campus-wide network for a new manufacturing facility. The proposed network consists of five carrier sense multiple access with collision detection (CSMA/CD) networks on a fiber distributed data interface (FDDI) backbone. In Section 1 the system configuration, the projected traffic pattern, and the proposed protocols are presented. Section 2 describes the models used in constructing the network simulation, while Section 3 contains the results and an analysis of the simulations. Some conclusions are drawn in Section 4.

1.0 THE ADVANCED SOLID ROCKET MOTOR (ASRM) FACILITY

The Advanced Solid Rocket Motor (ASRM) facility at Yellow Creek near Corinth and Iuka, Mississippi is part of a National Aeronautics and Space Administration (NASA) program to develop new solid rocket motors for the Space Shuttle. The facility will be a government-owned, contractor-operated operation. Lockheed Missiles & Space Company, Inc., ASRM Division is the prime contractor and the operation of the facility will be directed by the subcontractor Aerojet ASRM Division. RUST International Corporation is responsible for the engineering and construction of the facility.

Case preparation, propellant mixing, core stripping, core preparation, and motor assembly will be performed at the Yellow Creek facility. The facilities will be automated using the latest technology. An objective of this project is to provide a safe paperless environment. Therefore, a Local Area Network (LAN) is necessary to carry control and managerial information required to run the facility.

The facility will consist of several buildings spread over a large area. Manufacturing will take place in three separate buildings on the Yellow Creek site. A LAN will be required to provide communication for the manufacturing process. The network should be reliable, secure, and provide enough bandwidth to carry all the information [1].

1.1 ASRM system configuration

The proposed network consists of five 10 Megabits per second (Mbps) CSMA/CD networks linked together with a FDDI backbone. Most nodes talk only to the Operations Information System (OIS); there is little peer-to-peer communication. The OIS will also be on the FDDI backbone. The links are listed below:

- Link 1 – Support and Storage
- Link 2 – Final Assembly
- Link 3 – Propellant, Nondestructive Evaluation (NDE), Miscellaneous
- Link 4 – Case Preparation
- Link 5 – Mix Cast
- Link 6 – Operations Information System (OIS)

This work was supported by the National Aeronautics and Space Administration under Grant NAG8-866.

1.1.1 Operations Information System (OIS) The Operations Information System (OIS) is to provide an efficient means to plan and control the manufacturing of solid rocket motors for the ASRM project. The OIS is also the link between the business functions and the manufacturing functions of the facility. The OIS is a VAX cluster consisting of two VAX 6500's and two VAX 4300's. Scheduling, shop floor control, and data collection will be performed by the OIS; these functions will be provided by commercial software packages [1].

1.1.2 Medium The transmission media will be fiber optics which requires a separate fiber for receiving and transmitting. In each of the CSMA/CD LANs, optical hubs will connect the nodes together. Each receiving port of the hub retransmits the signal to all transmitting ports. Fiber optics was chosen because of its immunity to electromagnetic interference (EMI) and also to allow the network to upgrade easily to FDDI in the future if the need arises. Even though optical fibers are more expensive than the standard copper medium of transmission (twisted pair or coax), using fiber optics in a LAN offers several distinct advantages; especially for the given environment [2]:

- Because they use light instead of electricity, fiber optic cables are free from electromagnetic interference, crosstalk, and other types of noise except that which is introduced into the system from the electronic interfaces to the network. This is especially useful for sites with high levels of EMI.
- Since they have a large bandwidth with little inherent loss, optical fibers can provide data rates up to around 100 Gigabits per second (Gbps) over 100 km. One reason for this is that the bandwidth is inversely proportional to the length, while with wire the bandwidth is inversely proportional to the length squared.
- Due to the fact that taps are difficult to place in the network, optical fibers are very secure from unwanted intrusion.
- Since they are physically small and lightweight, optical fibers aid installation and maintenance. An optical fiber is generally 1/6 the weight of an equivalent coaxial cable carrying the same amount of information.
- Because optical fibers currently propagate with very little attenuation, typically as low as 0.2 dB/km, repeaters are not necessary for distances under 100 km.
- Since optical fibers carry no electrical current, they are ideal in situations where a spark could set off volatile substances.

1.2 ASRM Traffic Analysis

All nodes will communicate with the OIS. There will be three main types of nodes [1]:

- Workstations
- Workcells
- Area Supervisory Computers (ASC)

There will be several workstations throughout the facility available to the shop floor managers for entering and retrieving data during manufacturing. There will also be terminals in workcells. Each workcell has a specific job such as a vapor degreaser or a pattern cutting station. Some areas will have Area Supervisory

Computers (ASC) to archive data from several smaller nodes during manufacturing. This information will then be transferred to the OIS. The ASC must store the information in the event of a failed link to the OIS. After the link is restored, information must then be uploaded to the OIS.

1.2.1 Automative control and user response time Since control of the manufacturing will be accomplished over the LAN, the network should be reliable and have redundant links in case a link is lost. Also, the user response time is important because the user will be getting instructions from the OIS. The delay of information on the LAN for a given load is therefore an important consideration.

1.2.2 Data collection Large amounts of data must be stored because of the critical nature of the solid rocket motors in the Space Shuttle program. It is crucial that none of this data is corrupted or lost. Therefore, the network should be reliable, robust, and have redundant links.

1.3 Protocols

The manufacturing data network proposed for the Yellow Creek site is a hybrid of FDDI and CSMA/CD. The connections from building to building will be CSMA/CD based. FDDI will link the five CSMA/CD links together for processing in the OIS computer. All of the fibers installed in the system will be FDDI compatible 62.5 micron fibers to easily migrate to a full FDDI system, if the need arises.

1.3.1 FDDI FDDI, or fiber distributed data interface, is a network standard developed by the American National Standard Institute (ANSI X3T9.5) that operates at 100 Mbps. FDDI was started as a high-speed network to provide packet data between processors and fast storage devices. Now, FDDI is often used as a high-speed, low-error rate backbone to interconnect slower LAN's like IEEE 802.3, 802.4, or 802.5. FDDI uses optical fibers for the communication medium in networks with radius greater than a few hundred feet and has a timed token media access protocol with a ring topology. For the transmitting devices, light emitting diodes (LEDs) are generally used. By using multi-mode optical fibers, links around 2 km are standard. By using single-mode optical fibers with laser diode transmitters the link distance can be extended up to 60 km. FDDI has several distinct advantages over other protocols [3]:

- Up to 1000 connections.
- Total fiber path length up to 200 km.
- Bit error rate (BER) less than $2.5E - 10$

FDDI uses a form of serial baseband transmission that combines both the data and the clock transmissions in a single bit stream. Because the clock information is transferred with the data, synchronization is accomplished with the recovery of the data.

FDDI can use Manchester encoding, like Ethernet, but normally FDDI uses 4b/5b with NRZI encoding. 4b/5b means that it uses combinations of five code bits to represent a symbol of four bits. NRZI is an edge-type code that is short for "NonReturn to Zero Invert on ones" – which in optical fibers deals with polarity transitions. Every polarity change results in a logical "0" (low) while no change in polarity results in a logical "1" (high). Manchester encoding, on the other hand, is a level-type code where a "zero" starts at logic low and makes a low to high transition in the middle of a clock cycle, and a "one" starts at logic high and makes a high to low transition in the middle of a clock cycle. The 4b/5b NRZI coding is chosen over the standard Manchester encoding system that Ethernet uses for two major reasons:

- The 4b/5b with NRZI encoding is more efficient, requiring cheaper components.
- Along with the frame formats, the 4b/5b with NRZI encoding allows easier detection and correction of errors.

The 4b/5b encoding scheme is more efficient than Manchester encoding in that it converts four data bits to five code bits, resulting in an 80% efficiency. This requires the optical components to operate at 125 Mbps in order to obtain the standard 100 Mbps required for FDDI. With the Manchester encoding scheme, there are two pulses per data bit resulting in a worst case condition of 50% efficiency. This would require 200 Mbps components for the system to run at the standard 100 Mbps.

1.3.2 CSMA/CD The Ethernet system was initially designed by XEROX and uses carrier sense multiple access with collision detection (CSMA/CD). Many different stations are connected to a common bus. If a station has data to send and the bus is silent, the station will try to transmit a packet of data and then wait for an acknowledgement (ACK) from the receiving station. Once the receiving station sends an ACK, the transmitting station will send another packet. If two stations try to transmit at the same time then the information will collide, at which point each station waits a random amount of time before trying to transmit again. If the information from a station collides again, then the station waits a longer time before trying to transmit. Each station has an exponential backoff algorithm so the more collisions the longer each station will wait and the bus will quiet down.

The nominal data rate for Ethernet is 10 Mbps. Each station is connected to the coax at regular intervals of 2.5 meters to reduce reflections. The maximum link length is 2.5 km with repeaters. A maximum of 1024 stations can be connected to one Ethernet segment. Normally, coax cable is used to interconnect the computers although optical fibers and twisted pair can be used now. The standard topology is a bus topology.

The IEEE 802.3 CSMA/CD standard sends data in variable size frames commonly called "packets" with a minimum spacing of 9.6 us. The frame construction consists of [4]:

- 64 bit preamble
- 48 bit destination address
- 48 bit source address
- 16 bit type field
- 46 to 1500 bytes data field
- 32 bit CRC error check field

The preamble provides synchronization and frame mark. The destination address contains the physical addresses of a particular station or a group of stations. The source address contains the physical address of the transmitting station. The type field is used by high-level network protocols. The data field contains the data being sent. The error check field consists of a cyclic redundancy check (CRC) which is generated by the transmitting station. The receiving station generates a CRC when it receives a packet and checks it against the received CRC. If they do not match, then the transmission was garbled and the receiving station will ask for the packet again. This continues until an ACK is received, at which time the transmitting station can send another packet. The IEEE 802.3 standard allows 15 re-tries before the station times out.

1.4 Research Objective

The objective of the research is to analyze the proposed network to determine its performance at different loads. The two evaluation parameters used to judge the network are throughput and delay. The throughput is the effective bit rate of the system. It does not include the overhead bits used by the protocol or the packets that had to be transmitted again. The delay in a LAN is judged by the delay per packet.

2.0 SIMULATION WITH BONeS

The commercial software package BONeS [5] – Block Oriented Network Simulator – was purchased from Comdisco Systems, Inc. and installed. The software is written in LISP and allows the

user to graphically piece together blocks to model various networks such as FDDI, CSMA/CD, and X.25. For each part of the model it generates C source code. During a simulation, it links the code together and creates an executable to do the simulation.

BONeS is an event driven simulation. Each event has to be triggered by a previous event called a "trigger". If a block is not "triggered" then there will be no output. Therefore, when building a model using the provided blocks, race conditions must be avoided. Parallel inputs should be avoided. Instead, blocks should be cascaded to prevent race conditions.

2.1 Models

Models of CSMA/CD nodes, FDDI nodes, and bridges are included in the BONeS library. Also, an example of a campus-wide network is included in version 1.5.1 of BONeS [5]. These models were used to simulate the ASRM network consisting of five CSMA/CD LAN's linked together with a FDDI backbone.

2.1.1 CSMA/CD Workstation model BONeS comes with a complete model for a CSMA/CD workstation which includes the carrier sense, collision detection, exponential backoff, attempt limit, slot time, and the interframe gap. The parameters were set to the standard IEEE 802.3 CSMA/CD standards and are listed below:

- Backoff limit = 10
- Attempt limit = 16
- Slot time = 5.12×10^{-5} seconds
- Interframe gap = 96 bits
- Transmission speed = 1×10^7 bits per second

The following parameters were also set:

- Mean packet length = 6000 bits
- Propagation delay of an Ethernet transceiver = 2.0×10^{-6} seconds

2.1.2 FDDI backbone BONeS comes with a complete model for a FDDI workstation. Six FDDI workstation models were used to model the FDDI backbone of the ASRM network. The parameters are listed below [5]:

- Capacity = 100 Mbps
- Target Token Rotation Time = .01 seconds
- Operational Target Rotation Time = .01 seconds
- Propagation Delay between nodes = 1.0×10^{-5} seconds
- Ring Latency = 6.006×10^{-5} seconds
- Synchronous Allocation = 0.0 seconds
- Synchronous Buffer Size = 0
- Asynchronous Buffer Size = 2000 elements

2.1.3 Optical Hub model The model of the optical hub passes all frames that are received on the receive fiber to all transmitting ports with delay. This delay is caused by the light-to-electrical and electrical-to-light conversion. The delay in the optical hub was set to one nanosecond. The optical hub model is shown in Figure 1. The optical hub model is shown in Figure 1. The optical hub model is shown in Figure 1.

2.1.4 Traffic Source model A model to send a set number of packets randomly at Poisson intervals was developed. Once triggered by the Poisson generator the traffic source model sends a set number of packets as fast as possible. Another packet is sent as soon as a packet is sent successfully. The traffic source model is shown in Figure 2. The parameters for the traffic source model




Figure 1 Optical hub model




Figure 2 Traffic source model




Figure 3 ASRM network model

are listed below:

- interarrival mean of blocks
- number of packets

The model was set to send one packet per trigger.

The traffic source model was developed to model a workstation sending a block of data such as a text or graphics screen. During one iteration of a simulation, the traffic source model sends a set number of packets at an average interarrival rate set by the user. The interarrival rate has a Poisson distribution because traffic on a LAN tends to have a Poisson distribution.

2.2 Simulations

The complete ASRM network was modeled using BONeS. A model for each of the five CSMA/CD LAN's was developed and linked together via models for bridges to the model of the FDDI backbone. The complete network is shown in Figure 3; Link 4 (Case Preparation) is broken out in Figure 4 to show an example link. All CSMA/CD nodes were set to the IEEE 802.3 standards. The OIS computer was modeled as a single CSMA/CD node. All nodes sent packets only to the OIS per the ASRM communication model. The OIS sent a packet by randomly picking a node out of the address table. There were 129 CSMA/CD nodes in the simulation.

The traffic intensity per node was varied from 40 kbps to 89 kbps at ten points and the throughput and mean delay per packet for each of the links was collected using the probes provided in BONeS. The traffic intensity was varied with an exponential function to show the knees of the curves. The simulation time per iteration was set to ten seconds. The actual computer time to do the simulation was approximately 40 hours on a Sun Sparcstation II GX with 16 MB of memory.

3.0 SIMULATION ANALYSIS AND RESULTS

Figure 5 is a plot of the Mean Delay per Packet versus Traffic Intensity and Figure 6 is a plot of the Throughput versus Traffic Intensity. Both were created using the Post Processor in BONeS. All six links were plotted on each plot for comparison.

3.1 Delay per Packet

The Mean Delay per Packet versus Traffic Intensity plot is shown in Figure 5. Notice that all the links show a knee at a particular traffic intensity. Beyond this traffic intensity, many nodes are not able to transmit because of the heavy traffic. The default delay per node is zero. Thus, the mean delay of each LAN decreases once the LAN is overloaded.

Case Preparation




Figure 4 Case preparation model

ASRM Network Delay per Packet

ASRM Network Delay per Packet




Figure 5 Mean delay per packet versus traffic intensity




Figure 6 Throughput versus traffic intensity

Table 1 Maximum delay per packet and throughput for each of the links.

Link number	Number of nodes	Maximum delay per packet (milli-seconds)	Traffic intensity per node at maximum delay per packet (kbps)	Throughput at maximum delay per packet (Mbps)
Link 1	30	10.6	62	1.896
Link 2	20	26.4	68	1.288
Link 3	20	10.5	68	1.342
Link 4	27	5.5	68	1.821
Link 5	30	10.0	62	1.896

The traffic intensity at which the maximum delay per packet occurred for each link was recorded from the plot shown in Figure 5. This information is summarized in Table 1. Notice that Link 1 and Link 5 saturate first at about 62 kbps per node. This is because Link 1 and Link 5 have the most number of nodes. Link 1 and Link 5 have 30 nodes.

3.2 Throughput

The Throughput versus Traffic Intensity plot is shown in Figure 6. The curve shows the throughput increasing linearly with the traffic intensity per node. But, the analysis of the mean delay per packet shows that the links are overloaded beyond a certain traffic intensity. What is happening in the simulation is that only one node is able to transmit and all others cannot. The throughput beyond this knee is only available for one node. Therefore, the throughput for each of the links is compared at the traffic intensity where the maximum delay per packet occurred. The throughputs ranged from 1.288 Mbps for Link 2 to 1.896 Mbps for Link 1 and Link 5. This is shown in Table 1.

4.0 CONCLUSIONS

The FDDI ring provided ample bandwidth as was expected. However, there is a possibility of overload for each link. The proposed network is overloaded at 62 kbps per node for links one and

five and is overloaded at 68 kbps per node for links two, three, and four. This is assuming a worst case condition where all nodes are trying to transmit at the same time. Notice that the sum of the throughputs adds to 8.243 Mbps. This is approaching the maximum throughput of 10 Mbps for the OIS CSMA/CD node.

The analysis of the ASRM network was simplified by using the commercial software BONeS. The software allows the user to graphically build a network. The ASRM simulation was built using several components from the BONeS library. BONeS also allows the user to build his own modules. A very detailed simulation can be accomplished with BONeS. However, this also means that building a simulation can be a complex task and the actual time to do the simulations can be very long.

At this point in time, it has been proposed to use data over voice links for much of the non-manufacturing communications. This is primarily a cost saving proposal. Discussion of this proposal is beyond the scope of this paper.

REFERENCES

[1] Operations Information System Computer System Specification for ASRM, Rust International Corporation, 1990, Contract 21-3727.

- [2] Marion, R. F. Jr., "Optical fibers in local area networks", Advances in Local Area Networks, The Institute of Electrical and Electronics Engineers, Inc., New York New York, 1987, pp. 224-243.
- [3] Jain, Raj, "Error Characteristics of Fiber Distributed Data Interface (FDDI)", IEEE Transactions on Communications, August 1990, Vol. 38, No. 8, pp. 1244-1252.
- [4] Madron, Thomas William, Local area networks: the next generation, John Wiley & Sons, Inc. New York, 1990, pp. 176-199.
- [5] BONeS LAN Model Library Guide, USA: Comdisco Systems, Inc., 1991, pp. 9-1 ~ 9-28.

SECTION 2

CSMA/CD OVER FIBER WITHOUT FDDI BACKBONE

2.0 CSMA/CD OVER FIBER WITHOUT FDDI BACKBONE

Several changes have been made in planning the ASRM/YC facility. Buildings were deleted and computer terminals were moved. Also, the FDDI backbone was deleted from the communications system. Instead of the FDDI backbone, a VAXcluster with six redundant CSMA/CD ports was procured. Therefore, in late summer the network was updated and simulated to show the performance of CSMA/CD over fiber without the FDDI backbone. This will be used to compare CSMA/CD and the hybrid network discussed in Section 3.

2.1 Simulation models

The same BONeS models used to simulate the first network were used to simulate the CSMA/CD over fiber network. The parameters were set to the standard IEEE 802.3 parameters as before. It was assumed that each port in the VAX cluster could handle a full speed of 10 Mbps CSMA/CD. Therefore, each port of the OIS was treated as a separate CSMA/CD node. Also, the delays of commercial transceivers and optical hubs were used in the simulation [1–3].

2.2 Simulations

Five separate simulations were run to simulate the five different links and are shown in Figures 2.1 through 2.5. All nodes send packets to the OIS. The OIS randomly picks one of the nodes in the link and sends a packet to it. Each node has an equal opportunity of being picked. For a worst case analysis, the packet size was set to the smallest possible size of 64 bytes. The smaller the packet size with respect to the propagation delay the more collisions occur [4]. The load of each of the networks was varied from 1 Mbps to 10 Mbps in eight steps. The simulation time was set to vary with the load so that approximately 1000 packets were sent during each iteration.

Besides the propagation delays of the transceivers and optical hubs, delays caused by the length of the fiber cable between the OIS and the links were modeled. The distance–delays were calculated by dividing the distance by the speed at which the light travels down the fiber (.67 times the speed of light).

Simulation assumptions and parameters:

- All nodes talk to the OIS.
- OIS randomly picks a node and sends a packet to it.
- Each node sends one packet per trigger with a Poisson distribution.
- Propagation delay of an IEEE 802.3 transceiver = 500 nanoseconds
- Propagation delay of an optical hub = 630 nanoseconds
- Backoff limit = 10
- Attempt limit = 16
- Slot time = 5.12×10^{-5} seconds

- Interframe gap = 96 bits
- Mean packet length = 64 bytes
- Transmission speed = 1×10^7 bits per second

2.3 Analysis and results

The throughput and mean delay per packet were collected during each iteration for each of the five links. The Mean Delay per Packet versus the Normalized Throughput was plotted for each of the links on one plot and is shown in Figure 2.6. As the throughput increases, the mean delay per packet increases. The delay per packet at a throughput of 9 Mbps from Figure 2.6 was used to calculate the delay to send one graphics page. A page of graphics was assumed to be 640 pixels by 480 lines with 16 colors. This is equal to 153.6 kilobytes of data to be transmitted. The number of packets required to send 153.6 kilobytes was calculated using 64 byte packets. The delay per graphics page was calculated by multiplying the number of packets by the delay per packet. The results are summarized in Table 2.1.

Table 2.1 Delay per graphics page

Link	Number of nodes	Link delay (μsec)	Maximum number of cascaded hubs	Delay per packet at 9 Mbps (msec)	Delay per graphics page (sec)
Storage and Support	34	1.82	3	.242	.581
Final Assembly	15	4.85	3	.463	1.11
Propellant	12	0.758	2	.294	.706
Case Preparation	27	1.52	5	.262	.629
Mix Cast	40	5.46	4	.413	.991

2.4 Conclusions for second network

The simulations show the mean delay per graphics page to be on the order of one second. The largest mean delay of 1.11 seconds per graphics page occurred in the Final Assembly link even though it has only 15 nodes. Notice that the long length of fiber (approximately 3200 feet) causes a large propagation delay when compared to the other links. Also, it has three cascaded optical hubs. The mean delay per graphics page of Mix Cast was .991 seconds. Mix Cast also has a large propagation delay due to the distance and four cascaded optical hubs.

The packet size in the simulations was set to the smallest packet size of 64 bytes which is a worst case condition. Larger packet sizes would cause the delays to decrease [4].

The optical hubs should be arranged to cascade as few as possible. In fact, one commercially available optical hub specifies that each additional hub in a path reduces that path by 180 meters (590 feet) [5].

References

- [1] CentreCOM Micro Transceiver Series specification sheet. Allied Telesis, Inc., 1991.
- [2] CentreCOM Micro 125/126 Series specification sheet. Allied Telesis, Inc., 1991.
- [3] CentreCOM Micro 5000 Series specification sheet. Allied Telesis, Inc., 1991.
- [4] T. A. Gonsalves, "Measured performance of the ethernet," Advances in Local Area Networks. New York: IEEE Press, 1987, pp. 383–410.
- [5] Allen-Bradley ORnet Fiber Optic Ethernet System specification sheet. Allen-Bradley, 1989.

Storage & Support




Figure 2.1 Storage and support




Figure 2.2 Final assembly

Propellant

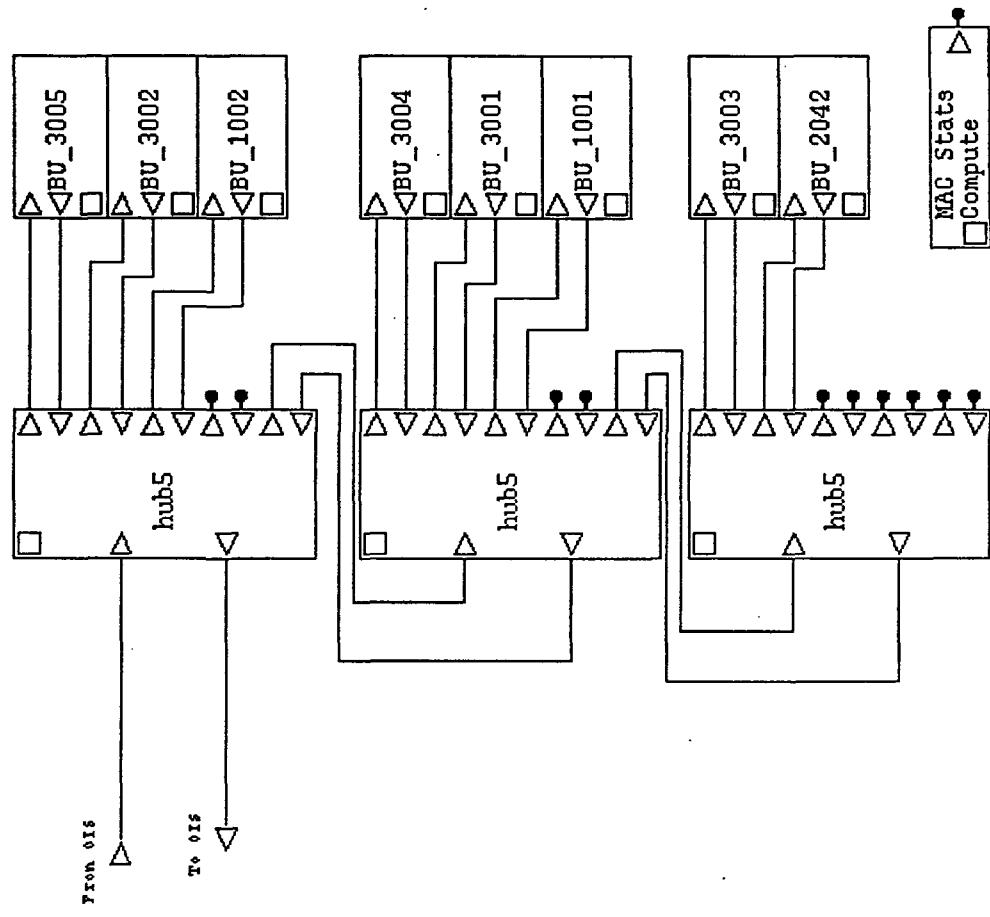


Figure 2.3 Propellant

Case Preparation




Figure 2.4 Case preparation




Figure 2.5 Mix cast




Figure 2.6 Mean delay per packet versus normalized throughput

SECTION 3

HYBRID NETWORK

3.0 HYBRID NETWORK

A second LAN technology was proposed by Intecom Inc. at a communications meeting in Iuka, Mississippi on July 9, 1991. Representatives from NASA, Lockheed, RUST, Mississippi State University, and Intecom attended that meeting. The proposed network, called LANmark, is a circuit switched network that integrates voice and data. It utilizes time slots to attempt to guarantee delivery.

A hybrid network consisting of CSMA/CD on fiber and LANmark is presently under consideration. Link 1 (Storage and Support) and Link 3 (Propellant) would be LANmark networks. Link 2 (Final Assembly) and Link 4 (Case Preparation) would be CSMA/CD on fiber. Link 5 (Mix Cast) would be part CSMA/CD on fiber and part LANmark.

3.1 LANmark

LANmark is a network which integrates voice and data switching [1–8]. A typical topology is shown in Figure 3.1. The LANmark Data Interfaces (LDI) encapsulate the IEEE 802.3 packet into a LANmark packet and send the packet to a LANmark Buffer located in the Interface Multiplexer (IM). The LANmark Buffer sends the packet to the central LANmark Packet Switch Board in the Integrated Business Exchange (IBX) which routes the packet to the correct IM. The IM then delivers the packet to the correct LDI.

3.1.1 LDI 400 The LDI 400 multiplexes voice and data. One telephone and one IEEE 802.3 data terminal can be connected to one LDI 400. The LDI 400 encapsulates the IEEE 802.3 packet into its own packet and sends it to the LANmark Buffer in the IM over twisted pair. The maximum distance from a LDI 400 to the IM is 2,000 feet. Sixteen LDI 400s can be connected to one LANmark Buffer Assembly.

3.1.2 LDI 410 The LDI 410 interfaces an Ethernet segment to the LANmark network. Since the LDI 410 is assumed to handle the aggregate load from multiple 802.3 sources, usually only one LDI 410 is connected to one Buffer Assembly.

3.1.3 LANmark Buffer Assembly The LANmark Buffer Assembly polls and buffers packets from the LDIs attached to it and sends them to the LANmark Packet Switch Board on fiber. The maximum distance from the IM to the LANmark Packet Switch Board is 100,000 feet on single-mode fiber. Up to fifteen LANmark Buffers can be installed in one Flex IM. The maximum bandwidth of each LANmark Buffer is 960 kilobits per second (kbps) full duplex. This bandwidth is available to the nodes connected to the LANmark Buffer on a polled basis. All nodes connected to a LANmark Buffer share the 960 kbps full duplex bandwidth. Figure 3.2 shows the arrangement of the LANmark Buffers and the LDI nodes [1].

3.1.4 Packet Board The heart of the LANmark system is the Switching Network (SN). An SN group (one card cage) that supports LANmark contains a SN Interface (SNIF), a Packet Bus Controller, and up to 8 SN controllers. Each SN controller that supports LANmark contains an IOB, a Packet Board, and an SN Processor Board. The Packet Board is the crucial link in the system. Our best reference on the Packet Board

[2] was provided by Glen Layfield, the Intecom salesman from Atlanta. Unfortunately, it is rather dated (March 1985) and we suspect out-of-date and somewhat inaccurate.

3.2 Nodes

Macintosh LC's have been procured as the workstations for the OIS network. They will use the X-Window system in communicating with the OIS. The load that this GUI puts on the network is only now being fully considered.

3.3 Analysis and results

The analysis of CSMA/CD over fiber is presented in Section 2 of the report. Calculations of the delay per graphics page were made using the measured delays per packet received from LANmark and are shown in Table 3.1. Also, a test of the performance of LANmark was performed on November 12–13, 1991 and on December 12–13, 1991 in Iuka, MS at the ASRM 791 building, Room 900. The results are presented in Sections 4 and 5.

Table 3.1 Delay of 153.6 kilobyte graphic page at varying packet sizes for LANmark

Packet size (bytes)	Delay of one packet (milliseconds)	Delay of 153.6 kbytes graphics page (seconds)
64	3.5	8.400
100	4	6.144
200	5	3.840
300	6	3.072
400	7	2.688
500	8	2.458
600	9	2.304
700	10	2.194
800	11	2.112
900	12	2.048
1000	13	1.997
1100	14	1.955
1200	15	1.920
1300	16	1.890
1400	17	1.865
1500	18	1.843

3.4 Conclusions for hybrid network

Assuming that only one node is talking, the LANmark system has an upper limit of approximately 2 Mbps full-duplex while the CSMA/CD has an upper limit of 10 Mbps. This causes longer delays for the LANmark system than for the CSMA/CD system. For

a single node, LANmark delays per graphic page are on the order of two seconds while the delays per graphic page are on the order of one second for CSMA/CD. See Sections 4 and 5 for the details of the LANmark tests.

References

- [1] Integrated Business Exchange, System Description, 590-2207-001, Issue 3, Intecom, Inc., 1989.
- [2] LANmark Packet Board, "Principles of Operation," ITCM-75-125, 590-2099-001, Issue 1, Intecom, Inc., 15 March 1985.
- [3] LANmark ETHERNET Support, "Operations and Maintenance," 590-2137-001, Issue 1, Intecom, Inc., 23 October 1987.
- [4] LANmark FLEX (LAN-Flex) Line Card Appendix, 590-2131-001, Issue 1, Intecom, Inc., July 1988
- [5] LANmark Packet Switch Functional Description, 070-1000-004, Rev. B0, Issue 2, Intecom, Inc., 4 October 1985.
- [6] LANmark Private Networks, User Guide, 590-2213-001, Intecom, Inc., October 1988.
- [7] LANmark 3270 Local Area Network, Engineering Applications, 590-2093-001, Issue 1, Intecom, Inc., 21 January 1985.
- [8] LANmark 3270 Configuration Guide, Applications Engineering, Applications Bulletin No. 8408, 590-2121-001, Issue 1, Intecom, Inc., 31 October 1985.




Figure 3.1 Typical LANmark topology




Figure 3.2 Arrangement of LANmark buffers and LDIs

SECTION 4
LANMARK TEST -- November 12-13, 1991

4.0 LANMARK TEST -- November 12-13, 1991

On Tuesday afternoon, November 12, 1991, the proposed network tests were commenced at the ASRM facility in Iuka, MS. The goal was to ascertain the relative performance differences of LANmark and fiber connectivity and to see if the Intecom/LANmark system could support DEC's LAT protocol. The participants included:

Michelle MacPherson (Lockheed)
John Donaldsen (Lockheed)
Robert Moorhead (Miss. State University)
Wayne Smith (Miss. State University)
Dale Scruggs (Rust)
Ron Wicks (DEC)
Mike Jones (DEC)
Eddie Orcutt (DEC)
Rich Drinkard (DEC)
Charlie King (DEC)
Rusty Lacy (DEC)
Merlin Hill (Aerojet)
Rod Wallace (Aerojet)
Jim Pulliam (Aerojet)
Travis Hawk (Aerojet)

The test configuration had changed from the configuration proposed at a meeting in Huntsville on 10/22/91. Michelle MacPherson had documented the supposed existing configuration as of 11/12/91. See Figure 4.1.

At a 3:00 p.m. meeting on Tuesday, November 12, it was decided to test the Intecom network to the best of our ability first. At 6:30 p.m. we commenced the tests by using Netcopy to transfer a large file from one of the VAX 6510s to one or more MACs. The file SYS\$SYS DEVICE:[000000]: INDEXFSYS;1 was arbitrarily chosen. It contained 2.26 Megabytes. Table 4.1 shows the transfer times. Note that for a few MACs the MAC ethernet port is the limiting factor, but as the number of data receivers (i.e., MACs) exceeds 3, the network becomes the limiting factor.

At this point, we decided to attempt to measure the response time of the network. A DECterm (X-window xterm) was initiated on every MAC. It took on the order of 5 minutes to get all the DECterms up; two in particular took two attempts. Rich Drinkard did a "show dev" on one MAC with no concomitant traffic on the other 14. It took 53 seconds for the show dev to show all the devices. In an attempt to see the effect of other traffic being on the network, the other 14 MACs then ran a "mon sys/int=1" while Rich ran another "show device". This time it took 59 seconds, an 11% slowdown. This didn't seem to be proving or showing much, so we tried to generate more traffic. With 14 MACs running "mon clus/int=1", a "show dev" command on the other machine took 67 seconds to complete. These results seemed inconclusive in deciding the suitability of Intecom/LANmark.




Figure 4.1: OSI Test Setup for the Intecom S/80

# of Macintoshes	machine #(s)	total time (sec)	total data rate supplied by network (kbits/sec)	data rate per MAC (kbits/sec)
1	12	45	400	400
1	12	45	400	400
2	12, 13	50	723	362
2	12, 13	56	646	323
3	12, 13, 14	68	798	266
3	12, 13, 14	66	822	274
6	{ 8, 9, 10 13, 14, 15	129	841	140
6	{ 10, 11, 12 13, 14, 15	129	841	140
15	1, . . . , 15	330	822	55

Table 4.1: DECCopy transfer rates over Intecom

At this point, DEC was interested in seeing if LAT would break. To attempt to do this, LAT sessions were run on 2 machines while the other 13 MACs generated different traffic pattern loads, such as simultaneously transferring the aforementioned file and running “show dev” and “mon clus/int=1” under DEC windows. The Intecom/LANmark network never really broke.

Rich Drinkard then created a file containing 80-column lines of the same characters, i.e.,

AA - - - - - AA
BB - - - - - BB
CC - - - - - CC
.
.
.
.
PP - - - - - PP
AA - - - - - AA
.
.
.
.

Scrolling through this file under LAT on three different machines first produced a loss of data (premature end-of-lines, dropped lines, etc.) However, this loss of data was due to buffer overflow. Turning on the Xon/Xoff protocol solved the problem. This proves that LAT in at least one case (in a straight ASCII file dump) does work over LANmark.

Finally, Rich's MAC which was running under LAT locked up. In an attempt to get the MAC communicating again he tried to reboot the MAC, then swapped MACs, and then

rebooted the DECserver. The DECserver would not reboot. The DEC personnel claim that the DECserver was communicating with the VAX (BIS? OIS?), but that it was timing out before it got the download. In an attempt to get the DECserver to reboot, Ron Wicks rebooted the LDI-400 to which it was attached. The LDI rebooted, but not the DECserver. DEC attributed the problem to the Intecom system.

We tried to get the HP Sniffer working, but we had no success. At midnite, we quit.

On Wednesday morning when the training area could not access the BIS (also known as the AIMS03 system), the LDI-400 on the DECserver was blamed because it was thought to be operating within the “Intecom Private Network”. To allow the DECserver to see the BIS, Aleisa Scott (Aerojet) added a LDI-400 port to which Rod Wallace connected an LDI-400 and connected the DECserver. This allowed the DECserver to see the BIS and for training to proceed. It now appears this may not have been the problem. Apparently the problem, at least to some degree, was that on at least some of the machines the LAT protocol connection had been selected. Apparently, this protocol would not allow access to the AIMS03 system and it was locking up the DECserver. Michelle MacPherson claims to have reconnected the old LDI-400 and brought up the AIMS03 system on the MACs, thereby showing that that LDI-400 is not working on a private network. There seemed and still seems to be a lot of confusion about what connections were made, when they were made, and how they were made.

514A	Flex Port	02. 000. 04	LDI - 410	(OIS connection)
781A	Flex Port	03. 000. 26	LDI - 400	
804B	Flex Port	03. 000. 09	LDI - 400	
808B	Flex Port	03. 002. 17	LDI - 400	
809A	Flex Port	03. 002. 18	LDI - 400	(old DECserver connection)
860A	Flex Port	03. 005. 24	LDI - 400	
816A	Flex Port	03. 003. 00	LDI - 400	
817B	Flex Port	03. 003. 03	LDI - 400	
821B	Flex Port	03. 003. 11	LDI - 400	
825A	Flex Port	03. 003. 18	LDI - 400	
829B	Flex Port	03. 003. 27	LDI - 400	
831B	Flex Port	03. 003. 31	LDI - 400	
836A	Flex Port	03. 004. 08	LDI - 400	
837A	Flex Port	03. 004. 10	LDI - 400	
838A	Flex Port	03. 004. 12	LDI - 400	
839A	Flex Port	03. 004. 14	LDI - 400	
840A	Flex Port	03. 004. 16	LDI - 400	

Table 4.2: IM. SLT. OffSet Port Card

On Wednesday morning, 11/13, the data in Table 4.2 was provided as the connectivity of Figure 1. The LDI-400 port to which the DECserver was now connected was not specified.

On Wednesday afternoon, 11/13, at 3:00 p.m. we reconvened the group to discuss what we had found out so far and to decide how to proceed. It was decided to rerun some of the tests from the previous night to verify we had the same setup and then to run multiple X-wave sessions on 14 MACs while trying various things on the other MAC to ascertain response time.

We started re-testing at 4:00 p.m. in hopes of having more net traffic. In re-running the tests we had 5 MACs download the same large file as the night before. It took 107 seconds, a transfer rate of 845 kbytes/sec. One MAC was being obstinate, but we were anxious to proceed, so we tried downloading from only 14 machines simultaneously. One machine had a session dropout, but the other 13 finished downloading in 303 seconds, a transfer rate of 835 kbytes/sec assuming 13 MACs. This, we felt, duplicated the performance of the previous night. The additional net traffic, if any, did not appear to have an effect.

We then ran X-wave demos on 10 MACs, while 5 MACs ran LAT sessions from the DECserver. (See Figure 4.2). LAT never bombed.

While most of us took a dinner break, Dale Scruggs ran 3 X-wave demos on every MAC and tried to transfer the 2.26 MByte file to one of the MACs running 3 copies of the X-wave demo. It took 23 minutes. This may say more about the MACs ethernet controller than anything else.

Also during this time, Merlin Hill set the HP Sniffer up on one of the LDI-400s (we therefore lost a MAC twisted-pair connection). This was almost useless since the Sniffer could only see broadcast packages and packages sent to it for the X-wave demo.

After supper, we ran 4 X-wave demos on 9 MACs, LAT sessions on 5 MACs and a paint package under X on one of the 9 MACs running X-wave. The response time was 2-10 seconds for paint commands to execute. Pop-up menus stayed up for 5 seconds and line drawing took up to 10 seconds. Even with all of this load, LAT never died while scrolling the aforementioned file of 80-column lines to the 5 MACs and 2 VT terminals (see Figure 4.2).

Travis Hawk checked the network statistics remotely and found the collision rate to be about 3% on the OIS ethernet while running 4 X-wave demos on 9 MACs plus the LAT load. The validity and cause of this number (3%) was later questioned and on November 20, 1991, it was reported by Lockheed (Michelle MacPherson) that at least some of our results were probably invalid. It was discovered that multiple connections existed from the VAX cluster in Figure 4.1 to the thick Ethernet and that these multiple connections were creating a loopback situation which was causing a large amount of collisions on the thick Ethernet.

The fiber connectors were mismatched, so we were unable to test a fiber network. We quit at 10:00 p.m.




Figure 4.2: Final Setup

SECTION 5
LANMARK TEST -- December 12-13, 1991

5.0 LANMARK TEST -- December 12-13, 1991

5.1 Introduction

The test group participants met at the Iuka ASRM facility at 3:00 p.m. on December 12, 1991 for the purpose of performing network testing on the communications network proposed for the Yellow Creek ASRM facility. An initial meeting took place at 3:00 p.m., and testing began at 6:00 p.m. The tests were concluded at 3:30 a.m. on December 13, 1991. The test plan, objectives, procedures and findings are presented in this report.

5.1.1 Test Participants

The test group consisted of personnel from the National Aeronautics and Space Administration (NASA), Lockheed Missiles and Space, Co. (LMSC), Aerojet (AAD), Rust Engineering (RUST), Digital Equipment Co. (DEC), Intecom, Inc. (INT) and Mississippi State University (MSU). The personnel in attendance (in no particular order) were:

John Donaldson (LSMC, Test Director)	
Merlin Hill (AAD)	Jim Pulliam (AAD)
Walter Robinson (NASA)	Charlie King (DEC)
Tom Kaeding (LMSC)	Robert Moorhead (MSU)
Wayne Smith (MSU)	Nancy Adams (DEC)
Mike Jones (DEC)	Eddie Orcutt (DEC)
Michelle MacPherson (LMSC)	Tom Reed (RUST)
Scottie Parma (RUST)	Pete Brewer (INT)
Dale Peterson (INT)	Mike Myrin (LMSC)
Ron Wicks (DEC)	Cliff Harris (LMSC)
Chris Allen (AAD)	Travis Hawk (AAD)
Phil Kelley (LMSC)	

5.1.2 Test Purpose

Five goals were stated as objectives for this OIS network configuration test. These goals were:

- A. Determine if the Intecom system as proposed for the Yellow Creek production environment is adequately and properly configured as specified in the Yellow Creek 60% Design.
- B. Determine if LAT, a time-sensitive protocol, will fail when utilizing the Intecom Packet Switch as the main Ethernet network carrier. Specifically, will an RS232 printer connected to a DECserver running LAT over the network fail due to a highly utilized Intecom network?
- C. Determine the relative difference in performance of the network traffic on the Intecom configuration versus the non-Intecom (fiber) configuration.
- D. Determine the relative difference in response time on the MACs communicating with the VAX via the network, where the network is configured with Intecom and without Intecom (fiber).

E. Determine, as best as possible with 15 MACintosh workstations, the maximum number of MACintosh X-Window workstations that can run on the non-Intecom network configuration.

5.1.3 Test Configuration

The major portion of the test was dedicated to testing the communications between the OIS thick Ethernet backbone that serves the OIS VAX cluster and a collection of 15 MACintosh workstations that were used to simulate the workload of a typical ASRM facility. In addition to the 15 workstations, a line printer was also connected to the backbone through a separate DECserver. Hence, the test essentially consisted of traffic between the 15 MAC workstations and the printer to/from the VAX cluster that represents the OIS network server.

Two different configurations were used for the test. One configuration (see Figure 5.1) was used for the non-Intecom (fiber) portion of the test. This configuration consisted of the following components:

DEC VAX 6510
Thick wire Ethernet backbone (731 Computer Room OIS)
DEC LANbridge 150 (NMC data collector)
DEC VAXstation 3100 Network Management Console (NMC)
DEC DEREN (Thick Ethernet to Fiber)
Fiber Optics Cable between buildings 731 and 791
Dec DEREN
Thick Ethernet Backbone (791 room 900)
Two DEC DELNIs
One DEC DECserver
One DEC printer (RS232)
Fifteen MAC workstations

The second test configuration replaced the fiber optics cable with the LANmark S/80 packet switch, and appropriate LDI devices to connect the switch to the thick Ethernet and the 15 MACs and the printer. This configuration (see Figure 5.2) was used for the Intecom portion of the test. This configuration consisted of the following components:

DEC VAX 6510
Thick wire Ethernet backbone (731 Computer Room OIS)
DEC LANbridge 150 (NMC data collector)
DEC VAXstation 3100 Network Management Console (NMC)
Intecom LDI 410 (Thick Ethernet to S/80)
Intecom S/80 packet switch
Fifteen Intecom LDI 400 (S/80 to MAC via twisted pair)
Fifteen MAC workstations
One LDI 400 (S/80 to DECserver)
One DEC DECserver
One DEC printer (RS232)

5.1.4 Test Sequence

The test plan consisted of two different network trials, one of which used the fiber network link and one which used the LANmark S/80 switch as a link. Within each trial, there were to be two subtests. The first subtest in each category was intended to develop a series of "load lines" indicating the load on the network as increasing numbers of MACintoshes were brought onto the network with increasing network demands. This test was to start with a single MACintosh running a single X-Wave application over the network. During a five minute period under this load, the average and peak network loads were to be recorded, as well as the average VAX CPU utilization.

After the first test, the network monitor counters were to be reset, and the test repeated with two MACs. After another five minute test with two MACs, the process was to be repeated with four, eight and then fifteen MACs each running a single X-Wave application. Throughout this test, the network and VAX performance parameters were to be recorded electronically by the network monitor and manually by personnel observing the test.

A second part of this "load line" test was intended to add additional windows on each MAC running additional copies of X-Wave to further increase the load on the network. It was originally intended that up to three X-Wave applications would be initiated on each MAC. As discussed below, this additional loading proved to be unnecessary.

The second subtest to be run on each configuration consisted of initiating multiple file transfers from the VAX to (and or from) the MACs over the network. These file transfers were intended to place a relatively heavy load on the network. A multi-megaByte file was transferred from the VAX to the MACs, beginning with one file transfer to one MAC, then two, then four, then eight, then fifteen. As in the previous subtest, the network parameters were to be monitored during the file transfer to determine the network performance under the varying load. These measurements were collected when the file transfers were approximately 60% complete. In addition, when multiple files were transferred, the total time required to transfer all the files was to be recorded. That is, the length of time required to complete all the file transfers was recorded as an additional measure of network performance.

Both subtests were to be performed with the fiber optics link, then the network would be restructured to the LANmark configuration (Figure 5.2), and the two subtests would be repeated with that network setup.

In order to test the reliability of LAT under all these varying network conditions, a single print job was to be started from the VAX to an RS-232 printer. This print job was several hours long, and was intended to continue throughout all four test phases. During the fiber link test, this printer was to be driven from one of the ports on one of the DELNIs in room 900, via a DECserver. During the LANmark test, the printer was to be driven by a DECserver through an LDI 400 over twisted pair from the S/80.

Throughout the tests, performance was to be measured by monitoring the traffic on the thick wire Ethernet backbone (731 Computer Room OIS) using a DEC LANbridge 150 (NMC data collector). During the fiber portions of the test, this monitor was connected

to the thick Ethernet backbone (791 room 900). During the LANmark portion of the test, this NMC remained connected directly to the thick wire Ethernet backbone of the OIS system via the fiber optics cable that was not used as part of the network during the LANmark tests.

5.2 Test Process

The test procedure as outlined above was initiated just after 5:00 p.m. when the network facilities were turned over to the test team. The first actual network data transfers were started at approximately 6:00 p.m. after the initial non-LANmark configuration had been established and verified.

5.2.1 Fiber Tests

The initial fiber tests went as planned. These initial tests are labeled as test numbers 1 through 5 in Table 5.1. Network performance measures are reflected in this table for the single X-Wave application for 1, 2, 4, 8 and 15 MACs. Both network utilization and CPU utilization increased pretty much as expected during this subtest.

Tests 6 through 8 represent the situations where additional X-Wave applications were added to five, ten and then fifteen MACs. After all 30 X-Wave applications were running, it was apparent that the additional applications were having little or no effect on any of the measurable parameters (compare test 5 and test 8 in Table 5.1). For this reason, it was decided not to add any additional X-Wave windows, and in fact, the second X-Wave application on all machines was removed for the subsequent tests, and a single X-Wave application was used for the remainder of the test.

While the additional X-Wave load did not seem to affect the network, it did appear to have a detrimental affect on the MACs. Two MACs locked up during test 8, and to save time, test 9 was run with only 13 MACs on the network (for some reason, it took on the order of ten to fifteen minutes to restart a MAC that had locked up). An additional MAC locked up during the first file transfer test (test 9), and hence test 10 was run with only 12 MACs on the network. Tests 11 through 14 were file transfer tests, and were run with 14 MACs on the network.

The file transfer tests over the fiber network are tests 9 through 13. In all cases, the file transfers were from the VAX to the MAC(s). The single file transfer, test 9, was used as a benchmark, and took about one and one-half minutes to complete. Test 10 transferred two copies of the same file from the VAX to two different MACs. While the network utilization increased, the total file transfer time for both files was not appreciably different from the single file. Because only twelve MACs were running during test 10, test 11 repeated the two file transfer test with fourteen X-Wave windows running on fourteen MACs. The results were essentially the same, and so test 12 was initiated with eight file transfers and test 13 was started with fourteen file transfers. While network and CPU utilization increased in the later tests, the file transfer time remained essentially the same. This will be analyzed below.

Throughout the fiber test, the RS232 printer continued to operate flawlessly. From this it was determined that LAT was operating properly under all load conditions.

5.2.2 Configuration Change

After test 13 was completed, the network was reconfigured to that shown in Figure 5.2. This is the LANmark configuration. The switch-over was delayed for a couple of hours by a faulty connector in the building 731 computer room. This was detected and corrected.

After the network configuration change, some difficulty was encountered in restarting the RS232 print job. The problem was eventually diagnosed as caused by a "private network" on the S/80 switch that would not let the DECserver load its software from the OIS network. A temporary connection was established that permitted the DECserver to load its software from the BIS network, after which the temporary connection was severed, and the DECserver operated the printer from the OIS network as in the previous test.

5.2.3 LANmark Tests

After the network configuration change, the previous tests were re-run using the LANmark network. To avoid confusion, the LANmark tests are labeled "A" through "J" in Table 5.1. Little difficulty was encountered once the tests began.

During the fiber tests, several MACs locked up and had to be restarted. One MAC failed during the fiber tests, and could not be restarted. Hence, the maximum number of MAC workstations used during the LANmark tests was 14. It is not likely that this loss of one machine had any measurable effect on the test results.

Tests "A" through "E" were the "load line" tests for LANmark. The network utilization and CPU figures are different than those from the fiber test, but show no unexpected deviations or variations. Based on experiences from the fiber tests, the multiple windows X-Wave applications were not repeated for the LANmark test, except for test "F", in which five additional X-Wave applications were used just to verify that these would have no effect on the tests. For the remainder of the file transfer tests, each MAC was running a single X-Wave application.

Tests "F", through "J" repeat the file transfer tests from the fiber experiment. Tests "F", "G", "H" and "J" repeat tests 9, 10 (or 11), 12 and 13, respectively. Test "I" was a bi-directional file transfer with four file transfers from the VAX to four MACs and four simultaneous MAC to VAX transfers. This bi-directional test was undertaken to test that the LANmark system would indeed provide approximately a one megabit data rate in each direction through the switch.

During tests "H", two MACs locked up, and during test "I" one MAC did. These were restarted before continuing with the tests. As with the fiber test, however, the RS232 printer functioned without a hitch throughout the LANmark portion of the test. This seems to indicate that LAT can operate very well in the LANmark environment.

5.3 Test Results and Conclusions

The recorded results from all tests are recorded in Table 5.1. Most of the results are plotted in Figure 5.3 through 5.12. These results can be used to answer the questions posed in the introduction to this report. To wit:


- A. Based on these test results, it does appear that the Intecom system as proposed for the Yellow Creek production environment is ade-

quately and properly configured as specified in the Yellow Creek 60% Design.


- B. The LAT time-sensitive protocol did not fail when utilizing the Intecom Packet Switch as the main Ethernet network carrier. Specifically, an RS232 printer connected to a DECserver running LAT over the network did not fail due to a highly utilized Intecom network.
- C. The relative difference in performance of the network traffic on the Intercom configuration versus the non-Intecom (fiber) configuration is included in Table 5.1. Specifically, the network utilization figures for both average and peak utilization can be gathered from this data. Tests 1 through 5 can be compared with tests "A" through "F" to determine this relative difference in performance.
- D. The relative difference in response time on the MACs communicating with the VAX via the network, where the network is configured with Intecom and without Intecom (fiber) can be determined by comparing the file transfer times for tests 9 through 13 with those for tests "F" through "J", respectively (tests 11 and "I" excepted).
- E. From this set of tests, it has not been possible to determine with any degree of certainty the maximum number of MACintosh X-Window workstations that can run on the non-Intecom network configuration. Certainly, it is clear that the number is greater than fifteen. Additional tests would be required to make a definitive determination of this value.

In general, the tests did not reveal any startlingly new information. The 10 megabit fiber optic Ethernet outperformed the one megabit (or two megabit if we use the bi-directional data rates) LANmark S/80, but that was expected. The relative network utilization figures and relative response times will have to be evaluated in terms of anticipated data rates for the OIS network before a determination of cost effectiveness can be made. LAT appears to operate well in the Intecom environment.

791



731



NO CONNECTION TO DEC OFFICE AREA

Figure 5.1: OIS Test Setup for Fiber Optics Cable




Figure 5.2: OIS Test Setup for the Intecom S/80

Test	No	No	No	FT	Network Util	CPU	Collisions		
	No	MACs	X-Waves	FT	Time	Avg %	Peak %	%	Mult Sing Defr
Load test -- over fiber									
1	1	1	0		1.6	1.9	19	1	0
2	2	2	0		2.9	3.1	21	0	0
3	4	4	0		5.3	6.1	26	2	3
4	8	8	0		9.5	10.8	37	3	4
5	15	15	0		18.3	19.2	51	6	1
Stress test -- over fiber									
6	15	20	0		17	19.2	52	23	10
7	15	25	0		17	19.2	53	3	0
8	15	30	0		17.5	19.2	51	12	4
9	13	26	1	1:27	20.5	21.8	52	6	5
10	12	24	2	1:33	23.2	24	52	6	1
11	14	14	2	1:31	22.4	24	54	8	1
12	14	14	8	1:29	40.7	45.8	60	5	1
13	14	14	12	1:30	44.5	53.1	46	2	0
Load test -- over Lanmark									
A	1	1	0		1.3	2.5	18	0	1
B	2	2	0		2.3	3.3	18	0	0
C	4	4	0		3.9	5.4	24	1	0
D	8	8	0		7.7	9.4	27	3	0
E	14	14	0		9.3	9.9	25	2	0
Stress test -- over Lanmark									
F	14	19	1	5:35	9.3	10.7	28	7	2
G	14	14	2	4:49	9.3	10.7	29	4	2
H	14	14	8	9:23	9.3	10.7	27	2	2
I	14	14	4&4	7:37	14.2	16.3	21	3	4
J	14	14	12	11:47	9.3	10.7	21	?	?

Table 5.1: Test Results




Figure 5.3: Number of MACs vs. Average Network Use




Figure 5.4: Number of MACs vs. Peak Network Use




Figure 5.5: Number of MACs vs. VAX CPU Utilization




Figure 5.6: Number of X-Waves vs. Average Network Use




Figure 5.7: Number of X-Waves vs. Peak Network Use




Figure 5.8: Number of X-Waves vs. VAX CPU Utilization




Figure 5.9: Number of File Transfers vs. Average Network Use




Figure 5.10: Number of File Transfers vs. Peak Network Use




Figure 5.11: Number of File Transfers vs. VAX CPU Utilization




Figure 5.12: Number of File Transfers vs. Time

SECTION 6

SUMMARY AND CONCLUSIONS

An RFP for the ASRM Communications Networks is almost ready to be released and yet the communications requirements have yet to be finalized. The 60% design review document, a copy of which we received on or about November 20, 1991, is too vague. It also appears at this point that the design will probably be finalized by a network vendor (i.e., the successful bidder) and will probably be procured with far too great a weighting on lowest cost. The bottom line problem of the whole design remains:

No one has yet to come close to specifying the data transmission requirements at the Yellow Creek ASRM facility

Given that fact, how can you specify or procure one or more networks? The only data rates we have seen were extracted from RUST, Inc. on 25 September 1991 and are listed below:

<u>workcell</u>	<u>building</u>	<u>data</u>
117	1016	330k/12 hrs
118	1016	1132k/12 hrs
121	1016	444k/176 hrs
109	103 (Michoud)	1912k/24 hrs
141A	103 (Michoud)	160k/2 hrs
101B	103 (Michoud)	70k/ 36 hrs
112	103 (Michoud)	458k/ 35 hrs
161	1016	1016k/14 hrs
160	1016	2730k/14 hrs
128	103 (Michoud)	2856k/20 hrs
114S	103 (Michoud)	8k/2 hrs
120	103 (Michoud)	166k/2 hrs
111	103 (Michoud)	42k/24 hrs
106	103 (Michoud)	2252k/10 hrs
124	103 (Michoud)	27010k/4 hrs
169	1016	6,596k/20 hrs
141A	103 (Michoud)	27,010k/4 hrs

Note how minimal this data is! Workcell 169 generates only 44,000 bits/minute or less than 800 bits/sec. This is *4 orders of magnitude* less than the bandwidth of Ethernet! If this is all that has to be handled, LANmark should work well. If it is not (and we doubt very seriously it is), the data requirements need to be specified -- NOW.

As for Intecom/LANmark, the Packet Board is what "paces" the roundtrip communications (see Test Results in Sections 4 and 5). Basically what happens in the present test system is that the VAX sends a packet, the IM buffer fills, two Packet Boards in two separate SN controllers cooperate to move the packet from one IM to another, the LDI moves the packet from the IM buffer to the MAC, the MAC consumes the packet, sends an acknowledgement packet back, and the LANmark system waits until the path is

clear and sends the "ack" packet to the VAX, at which point the VAX sends another packet. As can be seen from the test results in Section 4 and 5, the Intecom/LANmark system is slow, but steady. Since all the applications tested utilized ack-ack protocols, neither end (the VAX or the MACs) got ahead of the network. Each only sent when asked to send. It appears that the Intecom/LANmark system is only usable for "dribbles" of data or slow file transfer. The LANmark system does not, however, seem to suffer from excessive loss of data. All the tests so far show it to be a reliable steady system, albeit slow. The response time in a heavily loaded environment (2-4 Mbits/sec) would be unacceptable to most users. Some other problems with Intecom include the extensive use of unshielded twisted-pair cabling and the fact that some of the solutions are presently promises for the near future (i.e., shielded twisted-pair cabling plan in the first quarter of 1992). Note that this is the company that provided us with 1985 documents when we asked for documentation in late 1991.

In the whole specification concept there seems to be a failure to take life cycle costs into account. Intecom's LANmark is a specialized implementation. It will take extra training to teach someone how to setup and maintain the system. ASRM will have to pay to train these people. It has been said that there is a second source for Intecom LANmark parts, yet nobody has stated any second sources. It seems to us that if you go with an Intecom LANmark network you are now at Intecom's mercy and Intecom has not been a very stable company in the last 5 years or so. Technically Intecom's LANmark appears to work solidly, but apparently fiscally they are not.

Another question that must be answered soon is whether all 6 links are going to be handled by one VAX? Can one VAX handle the load? Is one VAX going to be a dedicated network server? This is yet another reason to ascertain the data transmission requirements now.

In summary, the biggest problem is the persistent failure of anyone to specify real data transfer requirements. Why do you need optical fiber where you do? Why can you "get away with" an aggregate 1 Mbit/sec communication channel on the links where you suggest LANmark? Until we get this information we can not do the best job possible analyzing the potential network designs.

Appendix A

Meetings Attended

Dale Thompson attended a communications meeting in Iuka, MS, on July 9, 1991, at which Intecom presented their product(s). This meeting was a rescheduled meeting that was finalized at the last minute over the July 4 holiday period. Dr. Robert Moorhead was scheduled to attend the original meeting, but was unable to attend the rescheduled meeting.

Dr. Robert J. Moorhead attended a communications meeting in Huntsville, Alabama on October 22, 1991, to discuss the network.

Dr. Wayne D. Smith and Dale R. Thompson attended a conference on the simulation program BONeS --Block Oriented Network Simulator-- in Atlanta, Ga. on October 25, 1991. BONeS is the simulation program being used to simulate the network to be installed at the Yellow Creek ASRM facility.

Dr. Robert J. Moorhead and Dr. Wayne D. Smith participated in a test of the Intecom LANmark network at Iuka, MS, on November 12-13, 1991.

Dr. Robert J. Moorhead and Dr. Wayne D. Smith participated in a test of the fiber-optic network and the Intecom LANmark network at Iuka, MS, on December 12-13, 1991.